UCRL-JC-132745
PREPRINT

# 3D Simulations of Line Emission from ICF Capsules

H. Scott, S. Langer, M. M. Marinak

This paper was prepared for submittal to the
1998 Nuclear Explosives Development Conference
Las Vegas, NV
October 25-30, 1998

**October 1, 1998**

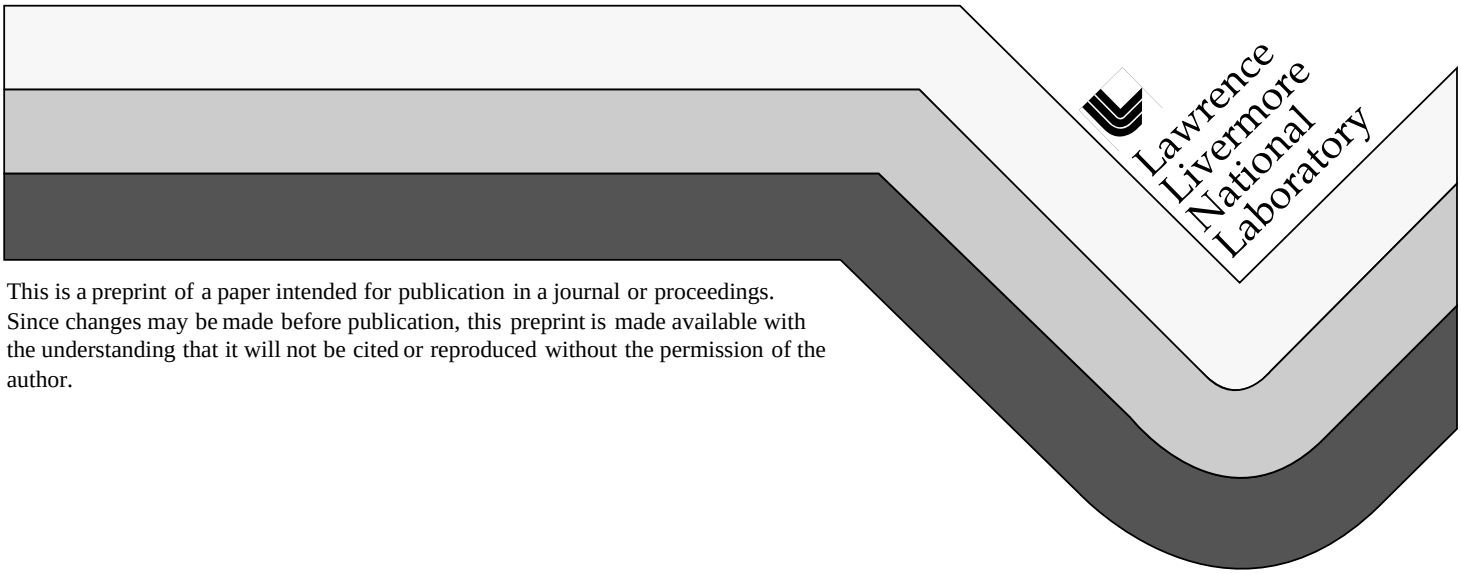

# 3D Simulations of Line Emission from ICF Capsules

**Howard Scott, Steven Langer, Michael M. Marinak**
**Lawrence Livermore National Laboratory**

*Line emission from ICF implosions can be used to diagnose the temperature of the DT fuel and provides an indication of the distortion in the fuel-pusher interface. 2D simulations have provided valuable insights into the usefulness of argon and titanium dopants as diagnostics of instabilities. Characterizing the effects of drive asymmetries requires 3D modeling with large demands for computer time and memory, necessitating the use of parallel computers. We present the results of some 3D simulations achieved with a code utilizing both shared memory and distributed parallelism. We discuss the code structure and related performance issues. (U)*

**Keywords:** ICF, Line radiation, parallelism

## Introduction

Drive asymmetries contribute to reduced yields in indirect drive implosions. Both power imbalances and pointing errors among the 10 beams of Nova or the 40 beams of Omega produce an asymmetric capsule and a distorted fuel-shell interface. Cold pusher material then pokes into the hot fuel and reduces yield. Rayleigh-Taylor (RT) growth of surface imperfections can also reduce yield in a similar manner.

2D simulations (Langer, et al., 1994) have provided valuable insights into the usefulness of argon and titanium dopants as diagnostics of instabilities. 3D models are needed to correctly model the saturation of RT modes and to model the drive asymmetry. Experimental techniques that distinguish between distorted fuel-shell interfaces and turbulent atomic mix at the interface are needed to understand the source of yield degradation. We hope to determine whether there are signatures of these two processes in the line emission.

## Summary of the Radiation-Hydrodynamic Model

The radiation-hydrodynamics simulations are done with HYDRA. We have done simulations for one Nova capsule and two Omega capsules, but present only the Nova capsule here. The capsule is depicted in Figure 1.

The simulation does not include a hohlraum and an external radiation source is used to drive the capsule. Marinak's mean opacity scheme is used so that only 3 radiation frequencies are needed. The capsule is driven by a 2.2 ns laser pulse with 32 kJ of 0.35 mm light. There is an intensity contrast of 5 between the foot and the drive pulse. The drive radiation used in the simulation is obtained from a hohlraum model and a view factor code, not a radiation-hydrodynamic simulation. The asymmetric drive radiation is applied directly to the capsule. The simulations cover one quadrant of the capsule and a run takes 1-2 weeks on 8 processors of a Cray J-90.

The x-ray drive asymmetry is typical of Nova experiments. The size of the asymmetry is based on observed pointing and power balance errors for the Nova laser. The time-dependent drive asymmetry on Nova includes power imbalance and beam mispointing. The errors in power balance and pointing are obtained from normal distributions which characterize the statistical averages for the Nova laser. The average beam-to-beam rms power imbalance during the foot is 10%, dropping to 5% at the peak. Pointing errors normal to the beam axis are 35 μm rms on

average. In this model, the laser power imbalance is constant throughout the pulse. The particular set of random numbers used to generate the drive in this simulation produced pointing errors larger than average (58 μm).

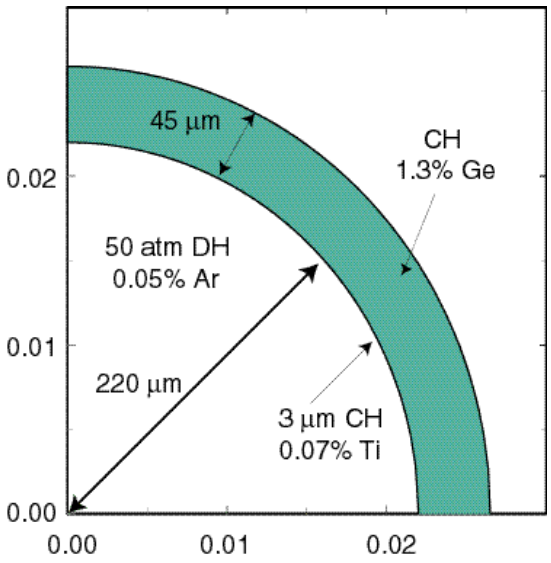


**Figure 1. The ICF capsule has a plastic shell filled with DH gas. The inner portion of the plastic shell is doped with titanium and the fuel is doped with argon. The plastic shell has a small amount of germanium to prevent higher energy drive X-rays from penetrating deep into the shell. The germanium was not included in the simulations discussed in this paper.**

The capsule implodes off-center due to a P1 component in the drive radiation. The material on one side of the shell becomes much denser than in the rest of the shell at peak compression. Figure 2 presents a density isosurface (for $\rho$=20 g/cm$^3$) for a time near peak burn. Line emission images (Figure 3) and upper state population densities (Figure 4) at this time also show off-center emission and capsule distortion.

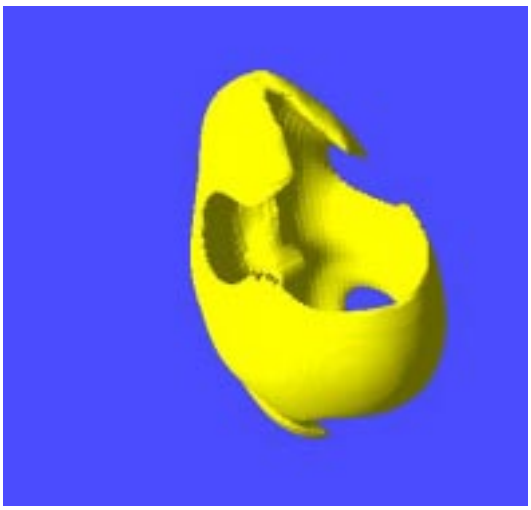

**Figure 2. 20 g/cc density iso-surface for the Nova capsule at time of peak burn.**

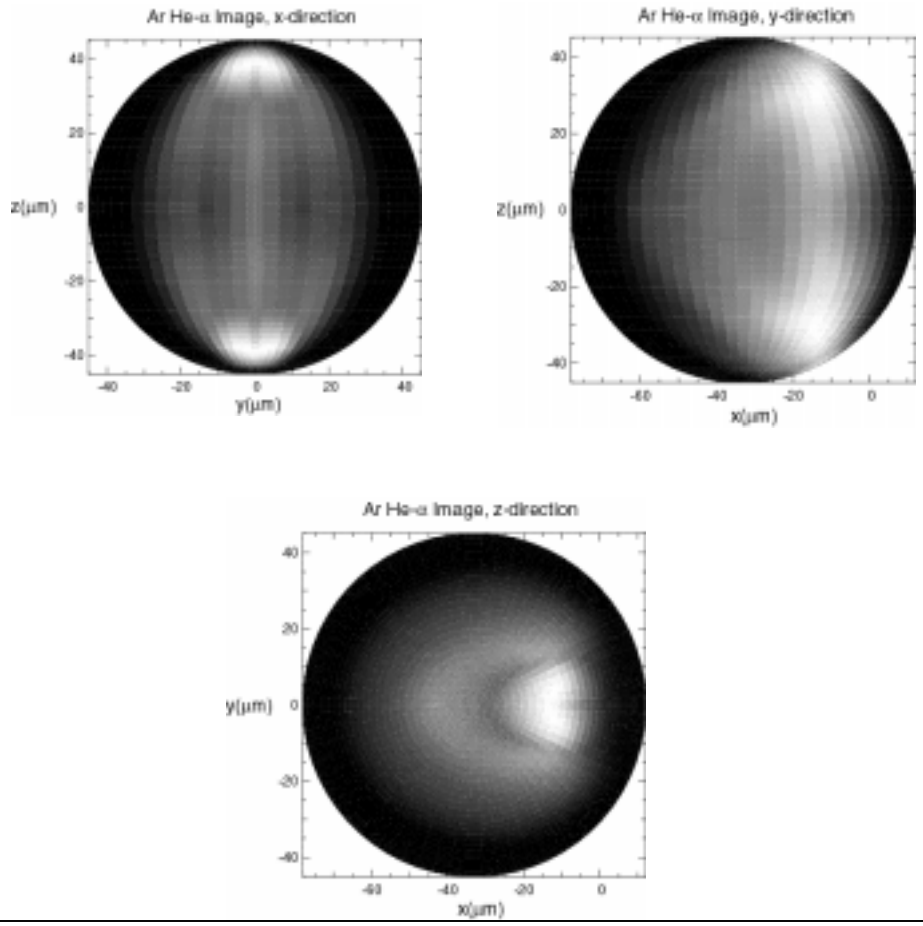


**Figure 3. Argon He-α images at peak burn in the x-, y-, and z-directions.**

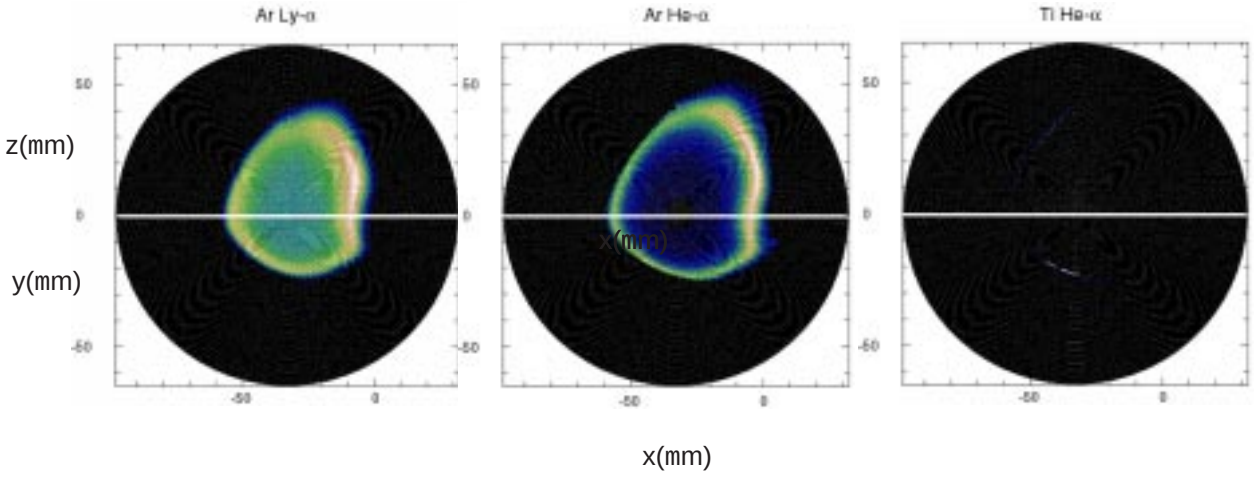


**Figure 4. Upper state population densities at peak burn for Argon Ly-α, Ar He-α and Titanium He-α lines.**

The intent of this paper is to discuss the computational aspects of the simulations, so we will not dwell on the implications of these results. For this aspect, we refer to Langer, et al. (1998).

**Summary of the line emission model**

The line emission is modeled with the non-LTE radiation transport code CRETIN (Scott and Mayle, 1994). Temperature and density profiles are extracted from the HYDRA simulation. CRETIN then self-consistently computes the populations of the atomic configurations, line radiation intensities and emergent spectrum.

The current CRETIN models use an atomic database from Ration and include only the argon in the DT fuel or the titanium in the shell, due to memory restrictions. The CRETIN simulations include only the innermost 120 radial zones, which contain all the argon and titanium. The simulations included only every 2nd or 3rd zone in $\theta$ and $\varphi$ from the HYDRA results, again due to memory limitations. This does resolve the largest amplitude modes at the time of peak burn.

**Parallelization strategies**

CRETIN utilizes both distributed-memory (DMP) and shared-memory parallelism (SMP). The distributed-memory parallelism uses MPI and the algorithms attempt to minimize communication. The shared-memory parallelism uses OpenMP directives. When both models are used simultaneously, only the master thread sends and receives messages. Both models of parallelism are implemented at a very high level. Each parallel region encompasses many subroutines. This decreases communication costs and increases the parallel efficiency. The trade-off for this is an increased sensitivity to load balancing.

Parallelization strategies differ for different physics packages. The atomic kinetics, i.e. the non-LTE calculations of atomic populations, opacities and emissivities, is parallelized over zones with both DMP and SMP. The zonal kinetics information (populations, opacities, etc.) is distributed over computational nodes, while the global information (mesh, atomic data, temperatures, mass densities, etc.) is replicated across nodes. The zones are coupled globally by radiation transport, which is split off into a different part of the simulation. Each zone is then completely independent of all other zones during this phase of the calculation, so there is no communication required during the atomic kinetics computations. Since the computational cost of zones can vary by several orders of magnitude, attention to load balancing is required. Currently, a static load balancing is accomplished by sorting zones by complexity before they are assigned to processors. The SMP dynamically load balances by sorting zones by cpu time on each node. Parallel efficiency and scaling have generally been very good with this strategy, although we would like to implement dynamical load balancing across nodes.

The radiation transport is treated as three different processes, according to the physics regime. The continuum radiation uses a moderate number of energies which are used in the atomic kinetics to calculate transition rates. The energies are essentially independent during the transport calculation (but are coupled during the atomic kinetics). In the DMP model, the transport calculations are parallelized over energies. The opacities, emissivities and intensities are distributed over nodes by energy. Since the atomic kinetics calculates (or uses) these quantities distributed by zone, this requires a global transpose of data both before and after the transport calculations, but no communication during the transport calculations. This scheme naturally load balances since the energies are computationally identical, but it doesn't scale to a large number of nodes, since the number of energies is moderate. The SMP implementation can either parallelize further over energies (on each node), or parallelize over directions. In the first

case, local work space can be shared between threads, while the second case requires additional work space for each thread. The amount of work in these calculations is generally a very small fraction of the total computational load.

The spectral radiation uses a large number of energies for the fine resolution of spectra, detailed line profiles, detector simulation, etc. It requires a separate calculation of opacities and emissivities from that performed by the atomic kinetics. The parallelization schemes are very similar to those used by the atomic kinetics and continuum radiation. For DMP, the opacity calculations are spread across zones and the transport is parallelized over energies. The SMP implementation further subdivides the work among threads, using a shared local work space. The communication issues are the same as for the combination of atomic kinetics and continuum radiation. Scaling and load balancing is naturally very good for this process, as the amount of work can be considerable for a high-resolution simulation.

The line radiation transport uses a moderate number energies to resolve line profiles. It treats small groups of lines simultaneously when physically necessary. A characteristic of line radiation is that the energies are very tightly coupled, as optical depths can change by orders of magnitude over the line profile. Groups of lines are treated simultaneously when the lines themselves are tight-coupled, either by virtue of significant frequency overlaps or through the atomic kinetics calculations. The DMP scheme respects these couplings by parallelizing over line groups. This minimizes the internodal communication, although it still requires a global transpose of (zonal) kinetic quantities. However, it may not scale to many nodes, depending on the application and the number of significant lines. The different line groups may have very different computational costs, but dynamic load balancing (via sorting by cpu time) can be used at negligible cost. The SMP scheme subdivides the work on each node by parallelizing over directions. The master thread handles the angle-independent portion of the calculations. This approach works well in terms of efficiency and load balancing, but requires that a significant amount of work space be duplicated for each thread. This has not been a significant issue for 2D calculations, but proved to be critical for these 3D calculations.

**Memory requirements**

The calculations are currently severely constrained by memory limitations. The HYDRA simulations were performed with a mesh size of 166 x 96 x64 zones. Including both the argon and titanium (as well as carbon and hydrogen) on a mesh with similar resolution results in the following parameters for the targeted CRETIN simulations: 149 x 98 x 66 nodes, 4 elements, 240 atomic levels, 10 lines, 15 energies/line, 21 continuum energies, 830 spectral energies, 128 angles. The aggregate memory requirement for the targeted simulations is ~27 GB.

For a parallel calculation, the memory requirements scale differently for each physics package and parallelization model. The DMP requirements for each package (for the targeted simulations) are summarized in Figure 5. Each package's requirements consists of a piece which is independent of the number of nodes plus a portion which scales inversely with the number of nodes. The atomic kinetics, continuum radiation and spectral radiation memory usage scale down to fit on current ASCI machines. The current line radiation algorithm does not scale well. It is suitable for small (i.e. 1D or 2D) problems. Less memory-intensive algorithms will require more communication but will be necessary for the calculations we would like to run

.

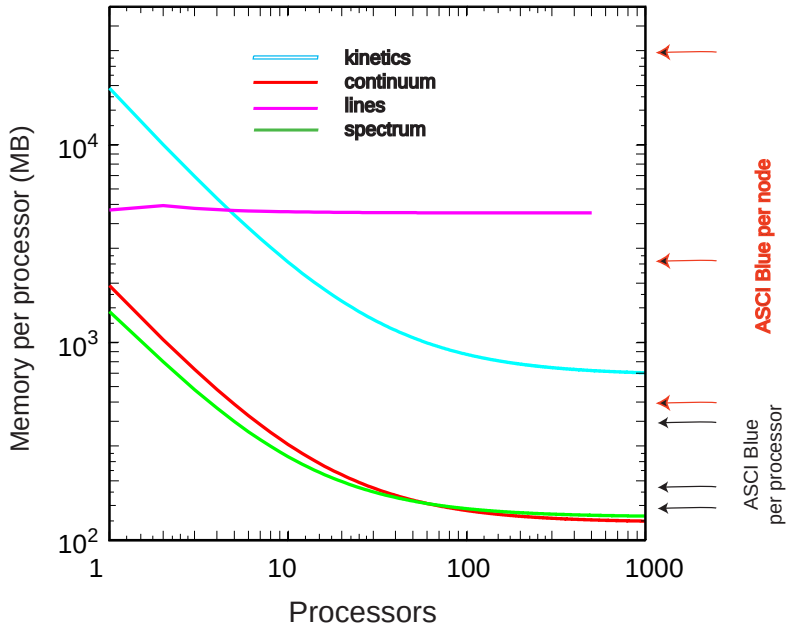


**Figure 5. Memory requirements per processor for targeted calculation for each physics package. Available memory on current ASCI Blue machines is shown on the right-hand side.**

Additional memory is required for each additional thread for both atomic kinetics and line radiation. For the targeted parameters, the atomic kinetics requirement is modest, amounting to only 1.4 MB per thread (depending primarily on the size of the atomic models). The line radiation memory requirement scales with the mesh size and number of line energies and amounts to 3.1 GB for each additional thread. Again, this is unsuitable for current machines.

The actual CRETIN simulations done so far used a mesh roughly 1/6 the targeted size. The parameters were: 120 x 34 x 34 nodes, 1 element, 70 atomic levels, 2 lines, and 282 spectral energies, resulting in an aggregate memory requirement of ~2 GB, which could be decreased to 0.6 GB per node.

**Progress Report**

The targeted calculations have not yet successfully run on the ASCI Blue machines, due to a combination of factors. The large memory requirements would make a realistic calculation problematic on the current version of Pacific Blue, an IBM SP2 with 168 nodes, each with 4 processors but only 0.5 GB memory. A memory upgrade (due shortly) would allow running the 1/6 size problem, but the IBM OpenMP compiler is not yet mature enough to compile the code correctly.

We also tried running these calculations on (unclassified) Mountain Blue, an SGI Origin 2000 system with 128 processors and 32 GB of memory per node. This system has adequate memory and the SGI OpenMP compiler does compile the code correctly (at least those sections required for these calculations), but does not execute correctly. The cause of this failure is currently under investigation.

We had a moderate amount of success with a new Sun system, which has 3 nodes, each with 16 processors and 3 GB of memory. The code compiled and executed correctly using KAI's Guide preprocessor for SMP mode. However, we found that we could only use 7 processors per

node in combined DMP / SMP mode, evidently due to some limitations in the operating system which are not well understood.

We obtained our current results (from the 1/6 size problem) on a single DEC Alpha (with 12 processors and 8 GB of memory), running in SMP mode only, compiled using KAI's Guide. The restriction to a single machine (using only a couple processors) was due only to the load on the cluster of Alpha's. The calculations took 11 hours per time step (total for all threads) and required many GB of disk space.

Future work will concentrate on decreasing the memory requirements of the line radiation transport. The current parallelization is adequate for 1D and 2D, but requires prohibitive amounts of memory for large 3D problems. Energy decomposition of the line radiation should be relatively straightforward to achieve and would decrease the memory requirements by roughly an order of magnitude. Spatial domain decomposition should be more effective in decreasing memory requirements but will require much more work to implement and it is uncertain how this will impact the effectiveness of the current algorithms. We are also considering using low-resolution calculations to provide effective escape factors to replace the line transport in high-resolution calculations. Performance scaling studies will follow.

## References

Langer, S. H., Scott, H. A., Keane, C. J., Landen, O. L., and Marinak, M. M., "Yield and Emission Line Ratios From ICF Target Implosions with Multi-Mode Rayleigh-Taylor Perturbations", *J. Quant. Spectrosc. Radiat. Transfer* **58**, 709-720 (1997).

Langer, S. H., Scott, H. A., Marinak, M. M., and Landen, O. L., "Modeling ICF Capsule Line Emission in Three Dimensions", to be published in *J. Quant. Spectrosc. Radiat. Transfer* (1998).

Scott, H. A. and Mayle, R. W., "GLF – A Simulation Code for X-Ray Lasers", *Appl. Phys.* **B58**, 35-43 (1994). The atomic kinetics and line radiation algorithms in CRETIN are very similar to those in GLF.



This work was performed under the auspices statement of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract No. W-7405-Eng.-48.